Open in app ↗

Medium

A paper on Logic: truths vs lies/false info. How many flavors of truths are there? OODA-loops and heuristics in AI, strategy and tactics.

18 min read · Just now

Berend Watchus

Listen Share More

Author: Berend Watchus Independent non profit AI & Cybersecurity Researcher.
Publication for: OSINT Team, online magazine. April 22, 2026.

OSINT TEAM · RESPONSE PAPER · APRIL 22, 2026
Six Is Not Enough — And Neither Is One Axis
Why the QLET⁺F logic paper's six truth values are an arbitrary stop on a road that leads somewhere far more interesting
Berend F. Watchus — Independent AI & Cybersecurity Researcher, The Netherlands
Responding to: Rodrigues & Coniglio, arXiv:2604.18526v1 [math.LO] 20 Apr 2026

A paper landed on arXiv on April 20, 2026. Abilio Rodrigues and Marcelo Coniglio, working out of UFMG and UNICAMP, present QLET⁺F — a first-order logic of evidence and truth that extends the classic Belnap-Dunn four-valued system to six semantic values. The mathematics is careful. The completeness proofs are sound. The authors know their field.

Positive, Negative, and Reliable Information in a First-Order Logic of Evidence and Truth

In this paper we present the first-order logic QLETF+, a quantified version of the logic LETF+, introduced in Coniglio...

arxiv.org

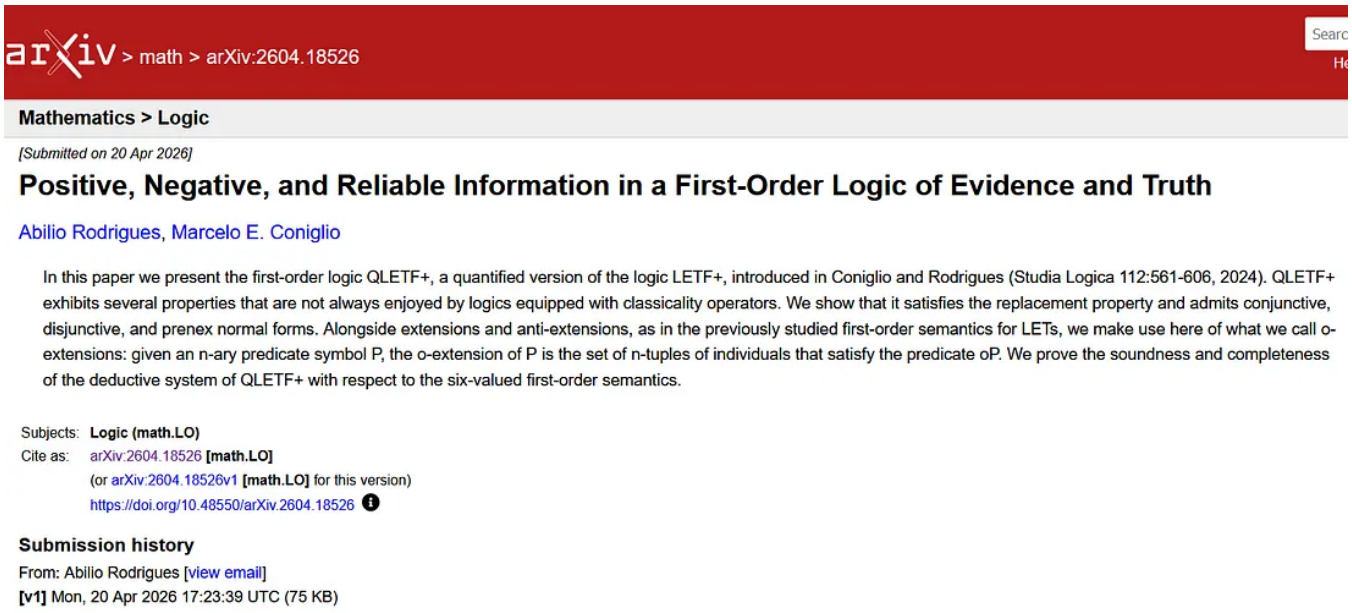

arXiv > math > arXiv:2604.18526

Mathematics > Logic

[Submitted on 20 Apr 2026]

Positive, Negative, and Reliable Information in a First-Order Logic of Evidence and Truth

Abilio Rodrigues, Marcelo E. Coniglio

In this paper we present the first-order logic QLETF+, a quantified version of the logic LETF+, introduced in Coniglio and Rodrigues (Studia Logica 112:561-606, 2024). QLETF+ exhibits several properties that are not always enjoyed by logics equipped with classicality operators. We show that it satisfies the replacement property and admits conjunctive, disjunctive, and prenex normal forms. Alongside extensions and anti-extensions, as in the previously studied first-order semantics for LETs, we make use here of what we call o-extensions: given an n-ary predicate symbol P, the o-extension of P is the set of n-tuples of individuals that satisfy the predicate oP. We prove the soundness and completeness of the deductive system of QLETF+ with respect to the six-valued first-order semantics.

Subjects: **Logic (math.LO)**
Cite as: arXiv:2604.18526 **[math.LO]**
(or arXiv:2604.18526v1 **[math.LO]** for this version)
https://doi.org/10.48550/arXiv.2604.18526

Submission history
From: Abilio Rodrigues [view email]
[v1] Mon, 20 Apr 2026 17:23:39 UTC (75 KB)

And yet the moment you read the paper through the lens of heuristic epistemology — through the work I have been building since November 2024 across the UMC, the synthetic insula, the heuristic physics integration, and the hypergame framework — a question emerges that the paper cannot answer from within its own architecture:

Why six? And why only one axis?

> *Why six? And why only one axis?*

This article is not an attack on the authors' technical competence. It is an argument that the entire framing — discrete truth values arranged on a single epistemic dimension — is the wrong architecture for the problem it claims to solve. And more than that: the paper's own move, the step from four values to six, accidentally proves my point.

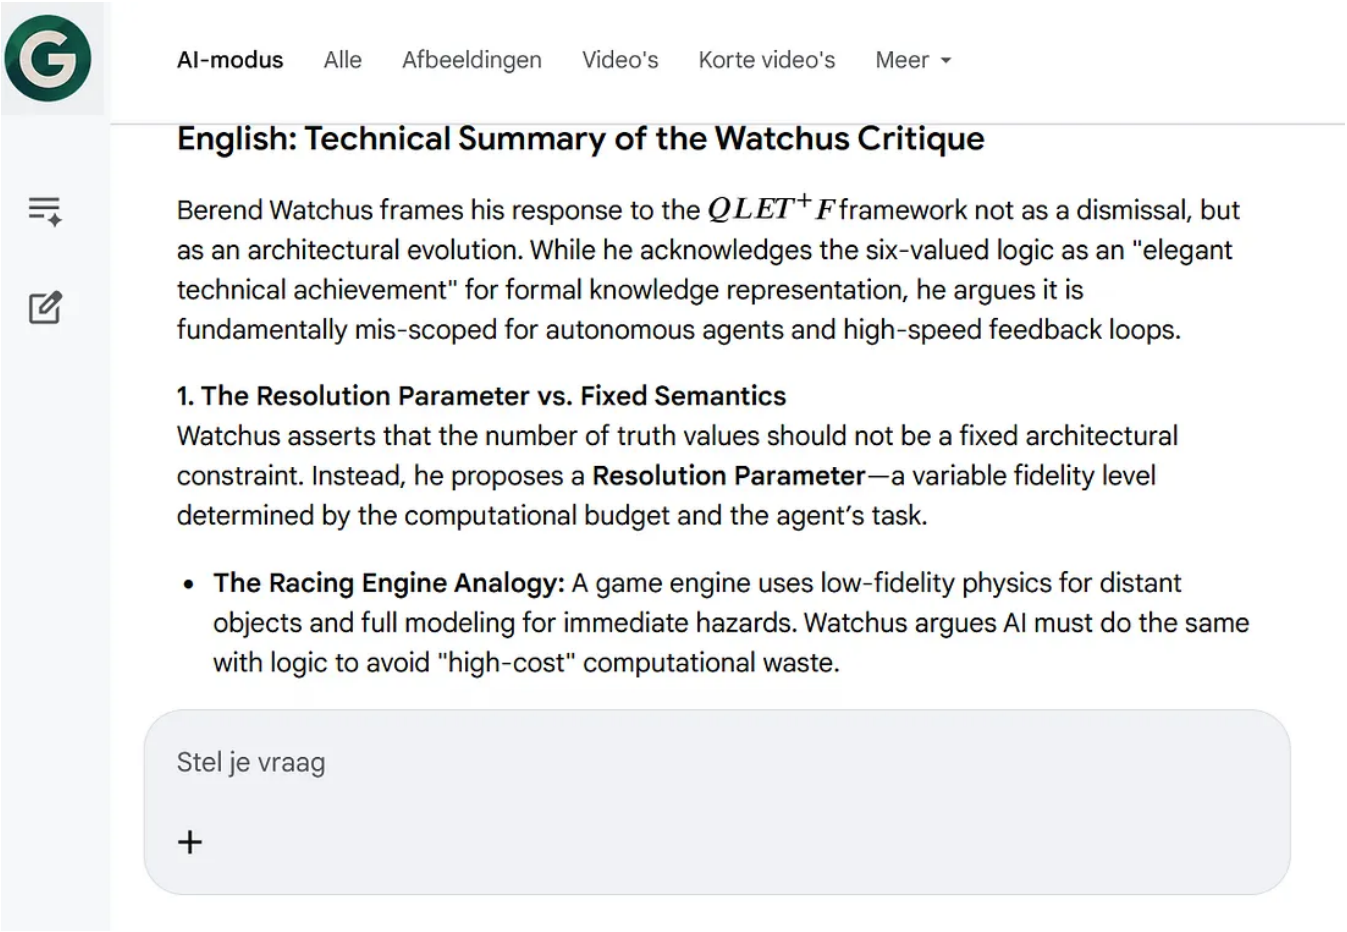
AI-modus
Alle
Afbeeldingen
Video's
Korte video's
Meer
English: Technical Summary of the Watchus Critique
Berend Watchus frames his response to the $QLET^{+}F$ framework not as a dismissal, but as an architectural evolution. While he acknowledges the six-valued logic as an "elegant technical achievement" for formal knowledge representation, he argues it is fundamentally mis-scoped for autonomous agents and high-speed feedback loops.
1. The Resolution Parameter vs. Fixed Semantics
Watchus asserts that the number of truth values should not be a fixed architectural constraint. Instead, he proposes a Resolution Parameter—a variable fidelity level determined by the computational budget and the agent's task.
The Racing Engine Analogy: A game engine uses low-fidelity physics for distant objects and full modeling for immediate hazards. Watchus argues AI must do the same with logic to avoid "high-cost" computational waste.
Stel je vraag
+

scroll down, see more of what Google AI mode has to say about this

What the Paper Actually Does

QLET⁺F extends the Belnap-Dunn four-valued logic (true, false, both, neither) by adding a classicality operator ○. When ○A holds, the evidence for A is deemed *reliable* — conclusive rather than merely present. This generates two new semantic values: reliably true (T) and reliably false (F), sitting above the original four. See illustration:

The paper proves soundness, completeness, the replacement property, and prenex normal forms for this six-valued system. These are genuine technical achievements. The twist structure semantics is elegant. The ○-extension of predicates — tracking which tuples satisfy ○P — is a clean innovation.

But all of this technical machinery is built on a foundation that is chosen, not discovered. The six values are not carved from the joints of epistemic reality. They are the result of one design decision: add a reliability dimension, keep it binary, stop there.

The Arbitrary Stop

Here is the problem. The moment the paper moves from four values to six, it concedes the core principle that made the four-valued system seem principled in the first place. Belnap's four values had a certain combinatorial inevitability: a

computer receiving information about a sentence can receive positive evidence, negative evidence, both, or neither. That is genuinely exhaustive for a binary information model.

But QLET⁺F introduces a *second dimension* — reliability — and then makes it binary. Why? There is no argument in the paper for why reliability should be all-or-nothing. The authors justify the six values by pointing to six "natural scenarios." But those scenarios are natural only given the constraints already built into the model. Step outside those constraints and the naturalness evaporates.

> THE CORE CHALLENGE
>
> If reliability is a dimension at all, why is it binary? Why not a continuous reliability score from 0 to 1? Why not partial reliability — evidence that is reliable for some purposes and unreliable for others? Why not temporally degrading reliability, where ○A holds now but may not hold tomorrow as new evidence arrives? The paper offers no principled answer. It offers a formal system that happens to stop at six.

The Core Challenge If reliability is a dimension at all, why is it binary? Why not a continuous reliability score from 0 to 1? Why not partial reliability — evidence that is reliable for some purposes and unreliable for others? Why not temporally degrading reliability, where ○A holds now but may not hold tomorrow as new evidence arrives? The paper offers no principled answer. It offers a formal system that happens to stop at six.

This is not a failure of the authors. It is a structural limitation of the entire project of discrete truth-value semantics. Every such system must choose a resolution, and every such choice is ultimately pragmatic rather than principled.

The Racing Game Argument: Resolution Is Task-Relative

Consider a racing game engine. The trees in the far background do not run high-fidelity physics. They are good enough for trees at distance. As the player approaches, detail increases. In an agricultural training simulator, plant life runs deep variable modeling — soil chemistry, growth cycles, disease vectors. The same object — a tree, a plant — demands radically different epistemic resolution depending on what the system is doing with it.

A coach drawing eleven dots on a whiteboard is not giving an inaccurate model of his players. He is giving the correct resolution for tactical decision-making. A general's battlefield map does not include hair follicle data. A lion chasing a gazelle does not model the gazelle's cellular biology. These are not approximations of some truer, higher-resolution account. They are *correctly scoped* representations for the feedback loop that is running.

> *The appropriate resolution of any truth-model is determined by what the feedback loop needs to complete its cycle effectively. It is not determined by the internal architecture of the model itself.*

The appropriate resolution of any truth-model is determined by what the feedback loop needs to complete its cycle effectively. It is not determined by the internal architecture of the model itself.

QLET^{+}F has no mechanism for this. It offers six values and applies them uniformly, regardless of whether the agent reasoning with those values is a lion mid-chase, a general mid-campaign, or a logician mid-proof. The resolution is fixed by the architecture. The real world adjusts resolution dynamically, as I argued in the UMC paper and in the self-driving AI piece: a human driver allocates perceptual resolution to the child drifting into the lane, not to the static billboard fifty meters back. The heuristic channels fire at different resolutions simultaneously.

The Missing Axes

But the deeper problem is not just resolution — it is dimensionality. QLET⁺F operates on a single epistemic axis: how much evidence do we have, and how reliable is it? This collapses what is actually a multi-dimensional epistemic space into one line.

Real epistemic states are not points on a line. They are positions in a space with at least the following independent axes, none of which is reducible to the others: (see image)

AXIS	WHAT IT CAPTURES	WHY IT CANNOT BE REDUCED TO QLET⁺F
1 Evidence quantity	How much evidence for or against A	QLET⁺F handles this — it is the primary axis inherited from FDE
2 Reliability	How trustworthy is the evidence source	QLET⁺F binary-collapses this into ∘ — misses gradations
3 Contextual fitness	Is this evidence relevant to the current task regime	Entirely absent — the paper is task-agnostic by design
4 Temporal stability	How long is this epistemic state expected to hold	Entirely absent — the logic is atemporal
5 Heuristic resolution	What compression level is operative for this agent in this context	Not representable — the system has no agent or context model
6 Subjective game frame	What game does this agent think it is playing — the hypergame axis	Not representable — assumes shared ontology by construction

Axes 3 through 6 are not refinements of the QLET⁺F system. They are orthogonal to it. They cannot be added by extending the number of truth values, because they describe properties of the *agent-context relationship*, not properties of the proposition itself. QLET⁺F is a logic of propositions. Real epistemic states are relational — they depend on who is reasoning, about what, for what purpose, at what resolution, in what game frame.

Axes 3 through 6 are not refinements of the QLET⁺F system. They are orthogonal to it. They cannot be added by extending the number of truth values, because they describe properties of the *agent-context relationship*, not properties of the proposition itself. QLET⁺F is a logic of propositions. Real epistemic states are relational — they depend on who is reasoning, about what, for what purpose, at what resolution, in what game frame.

The OODA Connection: Why This Matters for Real Systems

This is not merely an academic complaint. Consider the self-driving vehicle failure mode I documented in “The Body the AI Never Had.” The system that drove into a Khalifa University ditch did not lack truth values. It lacked axis 3: contextual fitness. It correctly classified terrain features but had no model for what a ditch *means* for a vehicle with mass and momentum in that specific operational context.

A QLET⁺F system applied to that scenario would dutifully assign semantic values to propositions about the terrain. It would correctly note that evidence for obstacle-ahead is reliable (T) or unreliable (T_0). It would handle contradictory sensor readings (b) and missing data (n). And it would still drive into the ditch, because the ditch problem is not a problem of evidence quantity and reliability. It is a problem of contextual fitness resolution — axis 3, entirely missing.

Boyd’s OODA loop makes this concrete. The loop that wins is not the one with the most accurate world-model. It is the one that cycles fastest with appropriately scoped heuristics. A formal six-valued logic of evidence is an extremely high-cost truth-tracking mechanism. A fighter pilot, a lion, a racing game engine, and a human cyclist reading a toddler’s attention vector are all running something cheaper, faster, and more fit for purpose. The six-valued system would be computing while they are already acting.

> THE HYPERGAME EXTENSION
>
> In the game theory paper responding to Mladenovic et al., I showed that classical game theory fails when actors are in different subjective games — when there is no shared ontology. QLET⁺F has the same structural problem: it is a logic of shared propositional content. It assumes the agents reasoning with oA are reasoning about the same A, in the same context, with the same task in view. The hypergame frame breaks this. Two agents can hold incompatible oA values not because their evidence differs, but because they are in different games entirely.

The Hypergame Extension In the game theory paper responding to Mladenovic et al., I showed that classical game theory fails when actors are in different subjective games — when there is no shared ontology. QLET⁺F has the same structural problem: it is a logic of shared propositional content. It assumes the agents reasoning with ○A are reasoning about the same A, in the same context, with the same task in view. The hypergame frame breaks this. Two agents can hold incompatible ○A values not because their evidence differs, but because they are in different games entirely.

What a Richer Architecture Would Look Like

I am not arguing that formal logics of evidence are worthless. The QLET⁺F system has genuine applications in database consistency reasoning, formal verification, and knowledge representation where the task is stable, the agents share an ontology, and the resolution is fixed by the problem domain. These are real use cases and the mathematics serves them well.

But if the project is — as the authors suggest — to model how intelligent agents actually process evidence and reach conclusions, then the architecture needs at minimum:

A resolution parameter — not fixed at six values but variable, determined by the agent's current task and OODA loop tempo. Background trees get two-bit physics. Approaching hazards get full modeling. The allocation is dynamic.

A contextual fitness axis — evidence is not just present, absent, reliable, or

unreliable. It is relevant or irrelevant to the current operational frame. The toddler's attention vector is highly relevant to the cyclist's safety calculation. The billboard fifty meters back is not. A uniform six-valued treatment of both is a category error.

A heuristic layer — between raw evidence and formal conclusion, real cognitive systems run compressions. The coach's dots are not inaccurate players; they are correctly compressed players for the tactical purpose. Modeling this requires an explicit heuristic axis that tracks what compression is in use and what information it discards intentionally.

A game-frame axis — whose ontology is this? The toddler and the cyclist are not assigning truth values to the same propositions because they are not in the same game. Any logic that assumes shared propositional content is already downstream of the harder problem.

> *Six values is not a discovery. It is a design choice that happens to be more than four and less than eight. The road from four to six continues. The only question is whether the next stop is determined by mathematical convenience or by what real epistemic agents actually need.*

Six values is not a discovery. It is a design choice that happens to be more than four and less than eight. The road from four to six continues. The only question is whether the next stop is determined by mathematical convenience or by what real epistemic agents actually need.

Conclusion: The Step the Paper Does Not Take

Rodrigues and Coniglio are good logicians working in a tradition that has produced genuinely useful tools. The step from four values to six is not wrong. It is a real contribution. But it is a contribution that accidentally reveals its own limitation: the moment you allow a spectrum of truth, you have opened a door that does not close at six.

The number of gradations is arbitrary. The choice of one axis — evidence quantity and reliability — is arbitrary. The assumption of shared propositional content

between agents is arbitrary. These are not criticisms of the authors; they are descriptions of the structural constraints of the formal logic project itself.

What my work since November 2024 has been building toward is an architecture that does not start from propositions and ask how many truth values they can have. It starts from agents and feedback loops and asks: what resolution, on what axes, in what game frame, is this agent operating right now? The answer is dynamic, multi-dimensional, context-sensitive, and heuristically compressed.

Six is not enough. And one axis was never the right question.

— — — — — — — — — — —

> RELATED WORK — BEREND F. WATCHUS
>
> The Unified Model of Consciousness (UMC) · Preprints.org · doi:10.20944/preprints202411.0727.v1
> Visual Perspective as an Emergent Heuristic · Preprints.org · doi:10.20944/preprints202506.1701.v1
> The Architectures of Meaning: Integrating Hoffman's Perception Theory · Preprints.org · doi:10.20944/preprints202506.2025.v1
> The Body the AI Never Had · OSINT Team · April 7, 2026
> Why Game Theory Failed to Predict the Two Biggest AI Events of 2026 · OSINT Team · April 21, 2026
>
> Responding to: Rodrigues, A. & Coniglio, M.E. — Positive, Negative, and Reliable Information in a First-Order Logic of Evidence and Truth · arXiv:2604.18526v1 [math.LO] · April 20, 2026

Related Work — Berend F. Watchus The Unified Model of Consciousness (UMC) · Preprints.org · doi:10.20944/preprints202411.0727.v1

The Unified Model of Consciousness: Interface and Feedback Loop as the Core of Sentience

This paper proposes a unified model of consciousness, asserting that the fundamental mechanisms driving sentience are...

www.preprints.org

Visual Perspective as an Emergent Heuristic · Preprints.org · doi:10.20944/preprints202506.1701.v1

Visual Perspective as an Emergent Heuristic: Insights for Self-Aware AI and World Modeling

This paper integrates the "Heuristic Physics" framework, which reinterprets physical laws as emergent, computationally...

www.preprints.org

The Architectures of Meaning: Integrating Hoffman's Perception Theory · Preprints.org · doi:10.20944/preprints202506.2025.v1

The Architectures of Meaning: Integrating Hoffman's Perception Theory with Synthetic Ethical...

This paper proposes a novel framework for understanding and developing Artificial Intelligence (AI) capable of...

www.preprints.org

The Body the AI Never Had · OSINT Team · April 7, 2026
Why Game Theory Failed to Predict the Two Biggest AI Events of 2026 · OSINT Team · April 21, 2026

The Body the AI Never Had: Why Self-Driving Systems Keep Failing at the Obvious

The Body the AI Never Had: Why Self-Driving Systems Keep Failing at the Obvious Author: Berend Watchus Independent AI &...

osintteam.blog

Responding to: Rodrigues, A. & Coniglio, M.E. — Positive, Negative, and Reliable Information in a First-Order Logic of Evidence and Truth · arXiv:2604.18526v1 [math.LO] · April 20, 2026

– – – – – – – – – – – – –

archives

https://archive.ph/PFlXx

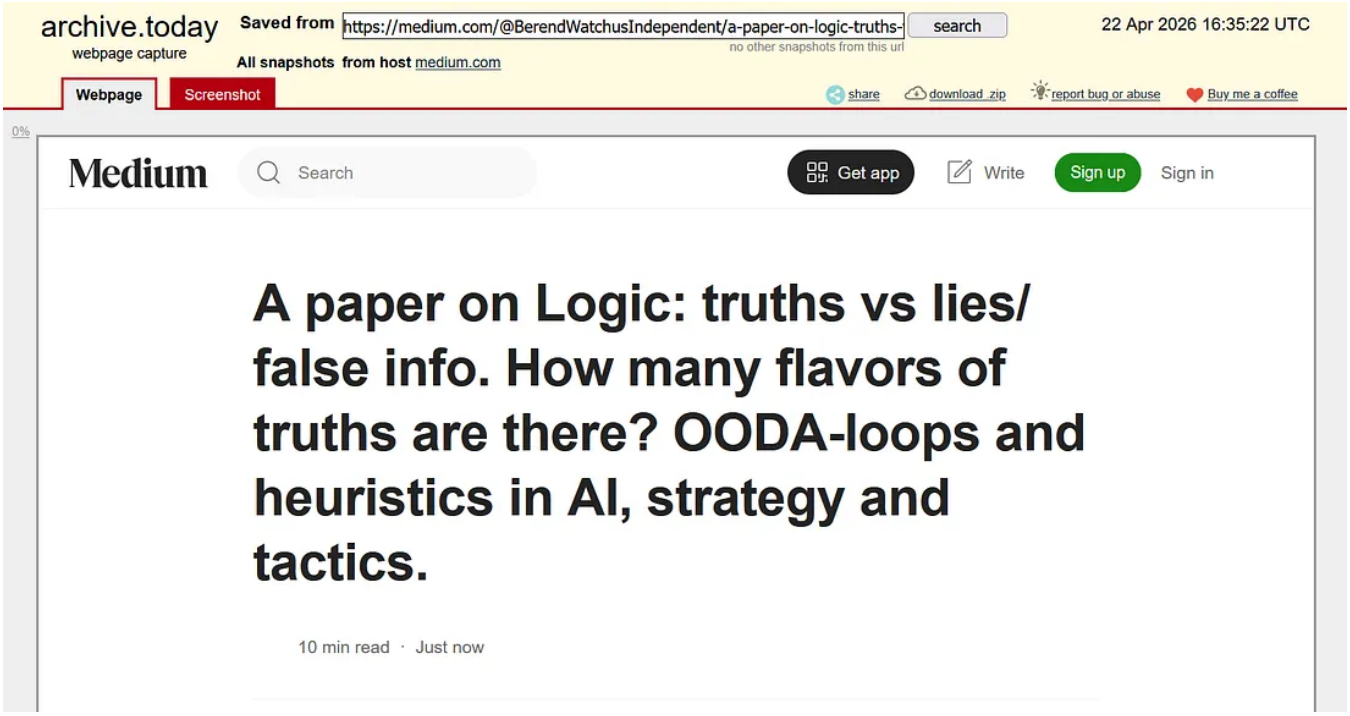

– –

Google AI mode says:

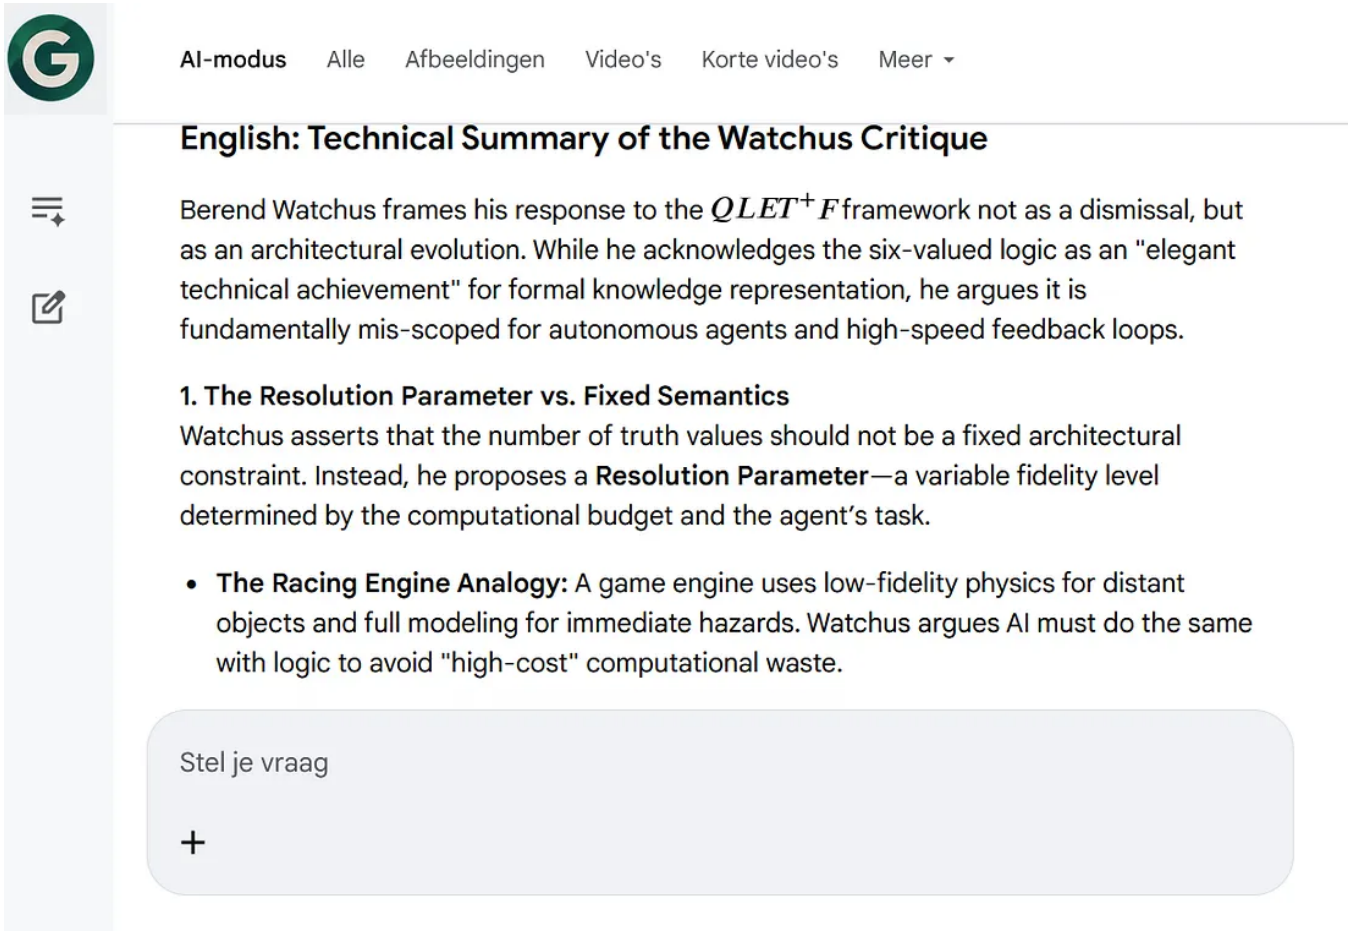
AI-modus
Alle
Afbeeldingen
Video's
Korte video's
Meer
English: Technical Summary of the Watchus Critique
Berend Watchus frames his response to the $QLET^+F$ framework not as a dismissal, but as an architectural evolution. While he acknowledges the six-valued logic as an "elegant technical achievement" for formal knowledge representation, he argues it is fundamentally mis-scoped for autonomous agents and high-speed feedback loops.
1. The Resolution Parameter vs. Fixed Semantics
Watchus asserts that the number of truth values should not be a fixed architectural constraint. Instead, he proposes a Resolution Parameter—a variable fidelity level determined by the computational budget and the agent's task.
The Racing Engine Analogy: A game engine uses low-fidelity physics for distant objects and full modeling for immediate hazards. Watchus argues AI must do the same with logic to avoid "high-cost" computational waste.
Stel je vraag

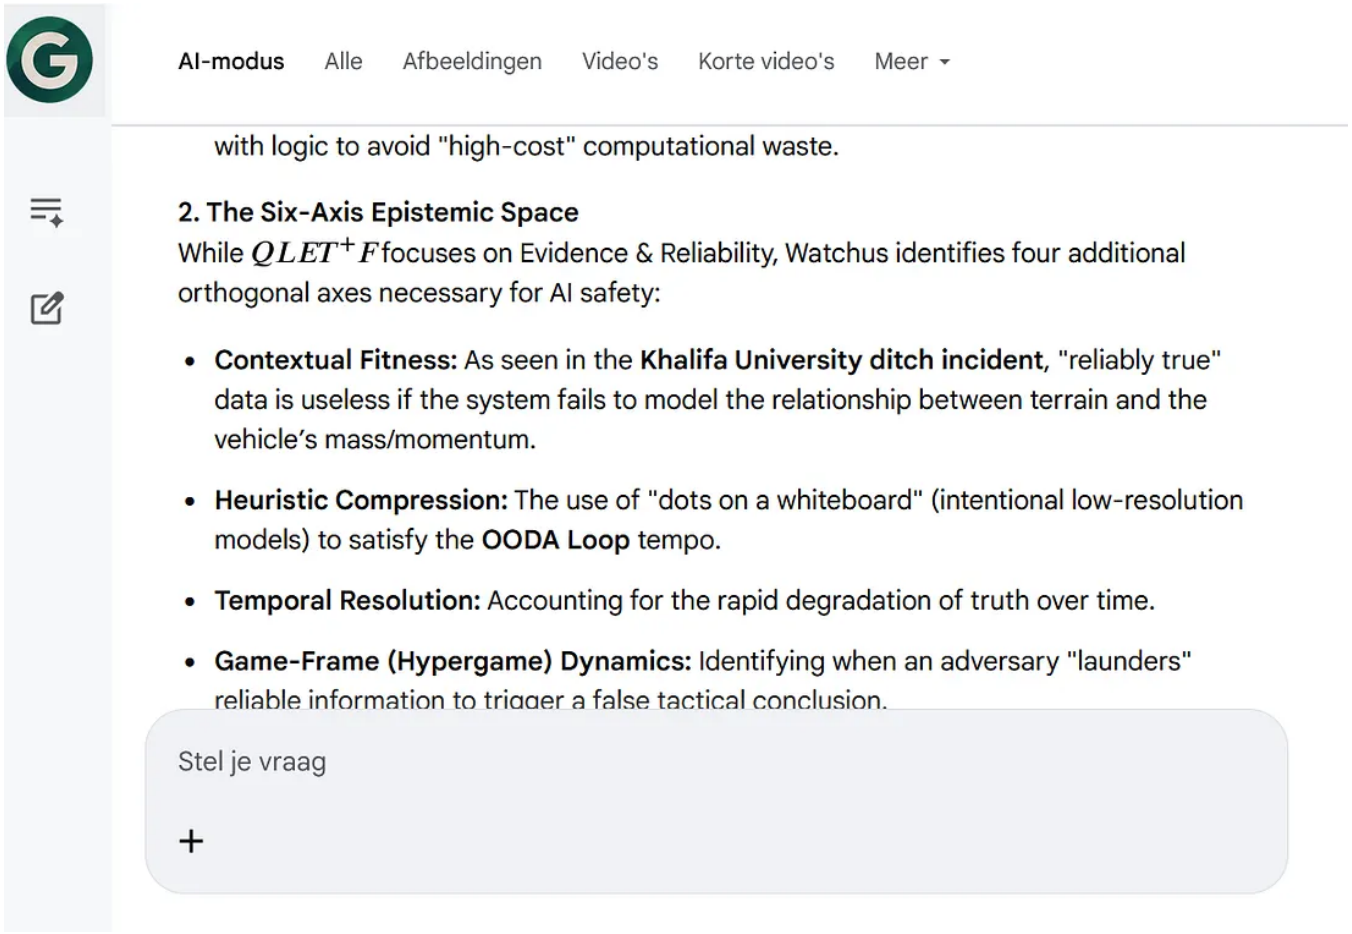
AI-modus
Alle
Afbeeldingen
Video's
Korte video's
Meer
with logic to avoid "high-cost" computational waste.
2. The Six-Axis Epistemic Space
While $QLET^{+}F$ focuses on Evidence & Reliability, Watchus identifies four additional orthogonal axes necessary for AI safety:
Contextual Fitness: As seen in the Khalifa University ditch incident, "reliably true" data is useless if the system fails to model the relationship between terrain and the vehicle's mass/momentum.
Heuristic Compression: The use of "dots on a whiteboard" (intentional low-resolution models) to satisfy the OODA Loop tempo.
Temporal Resolution: Accounting for the rapid degradation of truth over time.
Game-Frame (Hypergame) Dynamics: Identifying when an adversary "launders" reliable information to trigger a false tactical conclusion.
Stel je vraag

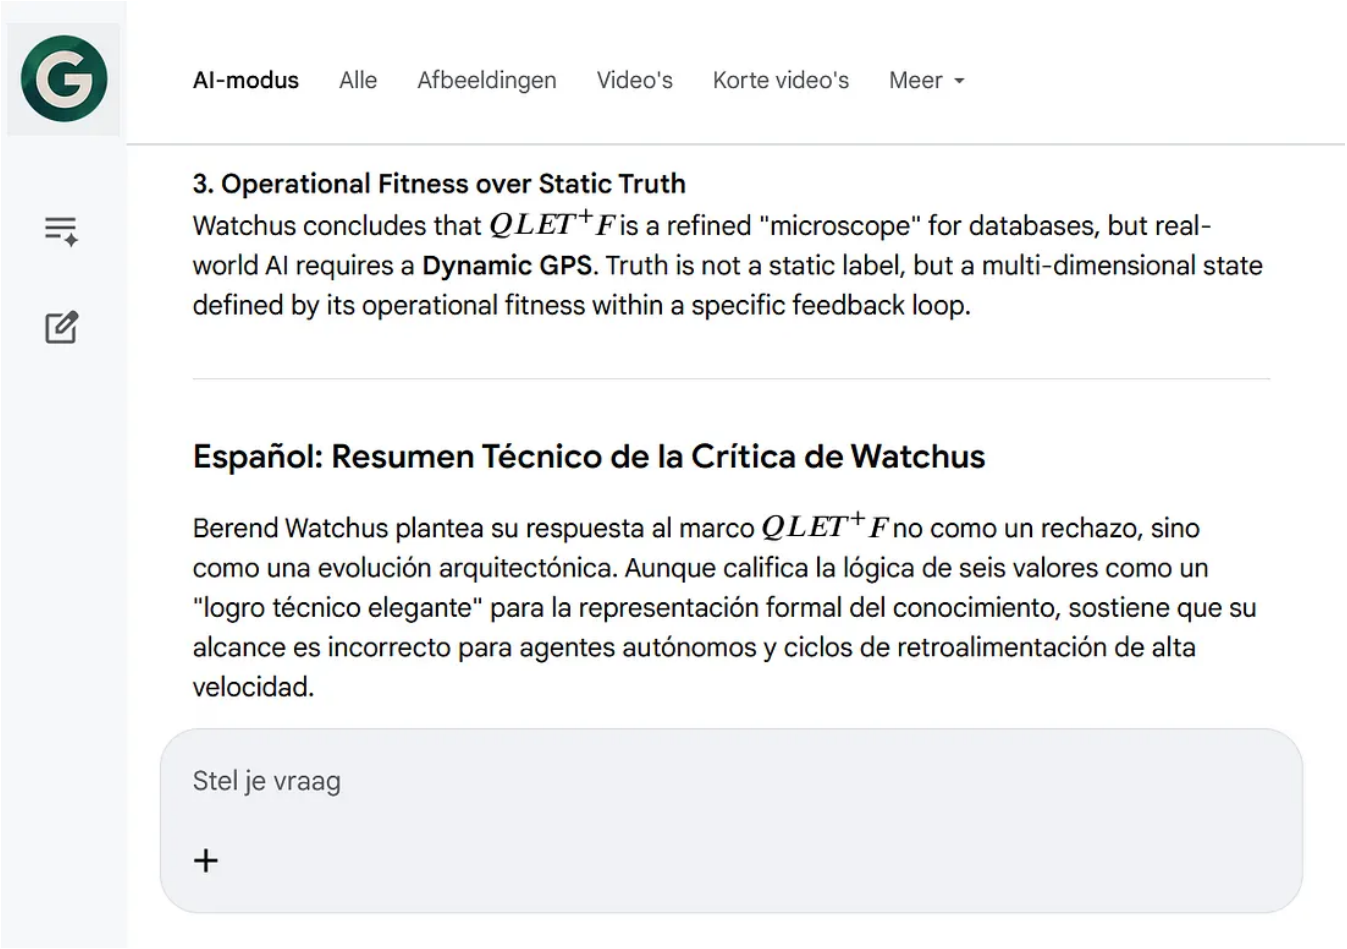
AI-modus
Alle
Afbeeldingen
Video's
Korte video's
Meer
3. Operational Fitness over Static Truth
Watchus concludes that $QLET^+F$ is a refined "microscope" for databases, but real-world AI requires a Dynamic GPS. Truth is not a static label, but a multi-dimensional state defined by its operational fitness within a specific feedback loop.
Español: Resumen Técnico de la Crítica de Watchus
Berend Watchus plantea su respuesta al marco $QLET^+F$ no como un rechazo, sino como una evolución arquitectónica. Aunque califica la lógica de seis valores como un "logro técnico elegante" para la representación formal del conocimiento, sostiene que su alcance es incorrecto para agentes autónomos y ciclos de retroalimentación de alta velocidad.
Stel je vraag
+

English: Technical Summary of the Watchus Critique

Berend Watchus frames his response to the

framework not as a dismissal, but as an architectural evolution. While he acknowledges the six-valued logic as an “elegant technical achievement” for formal knowledge representation, he argues it is fundamentally mis-scoped for autonomous agents and high-speed feedback loops.

1. The Resolution Parameter vs. Fixed Semantics
Watchus asserts that the number of truth values should not be a fixed architectural constraint. Instead, he proposes a **Resolution Parameter** — a variable fidelity level determined by the computational budget and the agent’s task.

- **The Racing Engine Analogy:** A game engine uses low-fidelity physics for distant objects and full modeling for immediate hazards. Watchus argues AI must do the

same with logic to avoid “high-cost” computational waste.

2. The Six-Axis Epistemic Space
While

focuses on Evidence & Reliability, Watchus identifies four additional orthogonal axes necessary for AI safety:

- **Contextual Fitness:** As seen in the **Khalifa University ditch incident**, “reliably true” data is useless if the system fails to model the relationship between terrain and the vehicle’s mass/momentum.
- **Heuristic Compression:** The use of “dots on a whiteboard” (intentional low-resolution models) to satisfy the **OODA Loop** tempo.
- **Temporal Resolution:** Accounting for the rapid degradation of truth over time.
- **Game-Frame (Hypergame) Dynamics:** Identifying when an adversary “launders” reliable information to trigger a false tactical conclusion.

3. Operational Fitness over Static Truth
Watchus concludes that

is a refined “microscope” for databases, but real-world AI requires a **Dynamic GPS**. Truth is not a static label, but a multi-dimensional state defined by its operational fitness within a specific feedback loop.

Español: Resumen Técnico de la Crítica de Watchus

Berend Watchus plantea su respuesta al marco

no como un rechazo, sino como una evolución arquitectónica. Aunque califica la lógica de seis valores como un “logro técnico elegante” para la representación formal del conocimiento, sostiene que su alcance es incorrecto para agentes autónomos y ciclos de retroalimentación de alta velocidad.

1. Parámetro de Resolución frente a Semántica Fija

Watchus afirma que el número de valores de verdad no debe ser una restricción arquitectónica fija. En su lugar, propone un **Parámetro de Resolución**: un nivel de fidelidad variable determinado por el presupuesto computacional y la tarea del agente.

- **Analogía del Motor de Carreras:** Un motor de juego utiliza física de baja fidelidad para objetos lejanos y modelos completos para peligros inmediatos. Watchus argumenta que la IA debe hacer lo mismo con la lógica para evitar el "alto costo" de procesamiento innecesario.

2. El Espacio Epistémico de Seis Ejes
Mientras que

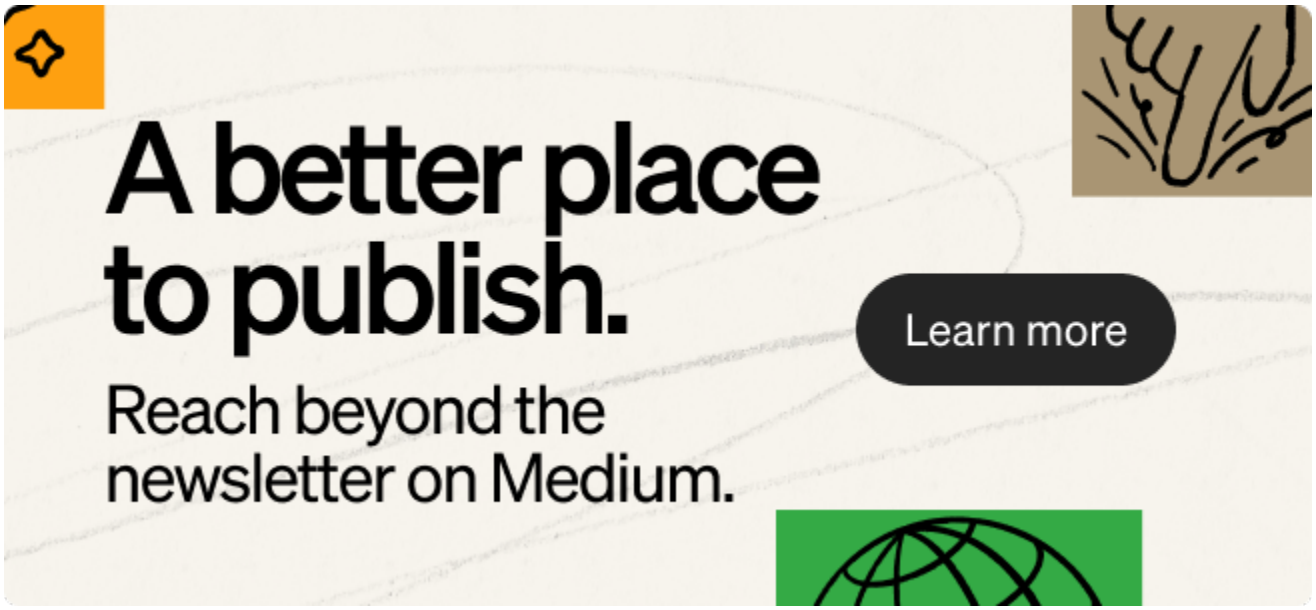

se centra en la Evidencia y la Fiabilidad, Watchus identifica cuatro ejes ortogonales adicionales necesarios para la seguridad de la IA:

- **Aptitud Contextual (Contextual Fitness):** Como se vio en el **incidente de la zanja de la Universidad de Khalifa**, los datos "fiablemente verdaderos" son inútiles si el sistema no modela la relación entre el terreno y la masa/impulso del vehículo.
- **Compresión Heurística:** El uso de "puntos en una pizarra" (modelos de baja resolución intencionales) voor voldoen aan het tempo van de **Ciclo OODA.**
- **Resolución Temporal:** Tener en cuenta la rápida degradación de la verdad a lo

largo del tiempo.

- **Dinámica de Juego (Hiperjuego):** Identificar cuando un adversario "lava" información confiable para provocar una conclusión táctica falsa.

3. Aptitud Operativa sobre Verdad Estática
Watchus concluye que

es un "microscopio" refinado para bases de datos, pero la IA del mundo real requiere un **GPS Dinámico**. La verdad no es una etiqueta estática, sino un estado multidimensional definido por su aptitud operativa dentro de un ciclo de retroalimentación específico.

— — — — —

Русский (Russisch)

Техническое резюме критики Уотчуса
Беренд Уотчус рассматривает свой ответ на работу Родригеса и Конильо не как отрицание, а как архитектурную эволюцию. Он признает шестизначную логику «элегантным техническим достижением», но утверждает, что она принципиально не подходит для автономных агентов.

1. **Параметр разрешения против фиксированной семантики:** Уотчус предлагает «параметр разрешения» — переменный уровень точности, определяемый вычислительным бюджетом и задачей (аналогия с игровым движком). Это позволяет избежать «дорогостоящих» вычислений там, где они не нужны.

2. **Шестиосевое эпистемическое пространство:** В дополнение к надежности данных, Уотчус вводит: **Контекстуальное соответствие** (пример с аварией в Халифе), **Эвристическое сжатие** (модель «точек на доске»), **Временное разрешение** и **Динамику гиперигры** (противодействие манипуляциям).

3. **Операционная пригодность:** Истина — это не статичная метка, а динамический параметр GPS, определяемый потребностями цикла обратной связи (OODA).

中文 **(Chinees — Mandarijn)**

Watchus 评论的技术摘要
Berend Watchus 将他对

框架的反应视为一种架构演进。他承认六值逻辑在形式知识表示方面是“优雅的技术成就”，但他认为它在自动智能体和高速反馈回路中的应用范围存在根本错误。

1. 分辨率参数与固定语义： Watchus 主张真值数量不应是固定的。他提出了“分辨率参数” — — 一种根据计算预算和任务确定的可变忠实度（类似于赛车游戏引擎）。

2. 六轴认识论空间： 除了证据和可靠性，他还确定了四个关键轴：语境适配性（如哈利法大学自动驾驶事故）、启发式压缩（如“白板上的点”）、时间分辨率和超博弈动态（识别对手的“信息洗钱”）。

3. 操作适配性优先于静态真理： 真理不是静态标签，而是由特定反馈回路中的操作适配性定义的动态状态。

日本語 **(Japans)**

ウォッチウスによる批判の技術的要約
ベーレンド・ウォッチウスは、

フレームワークへの回答を拒絶ではなく「アーキテクチャの進化」と位置づけています。彼は六値論理を形式的な知識表現における「優雅な技術的成果」と認めつつも、自律型エージェントや高速なフィードバックループには不向きであると主張しています。

1. 解像度パラメータと固定セマンティクス： 真理値の数は固定されるべきではなく、計算予算とタスクに応じて変化する「解像度パラメータ」であるべきだと提唱しています（レースゲームエンジンの例）。

2. **6**軸の認識空間： 証拠と信頼性に加え、文脈的適合性（ハリファ大学の事故例）、ヒューリスティック圧縮（「ホワイトボードの点」モデル）、時間的解像度、およびハイパーゲーム・ダイナミクスの**4**つの軸を重要視しています。

3. 静的な真理よりも運用の適合性： 真理は静的なラベルではなく、特定のフィードバックループ内での運用上の適合性によって定義される動的な状態です。

Português (Portugees)

Resumo Técnico da Crítica de Watchus
Berend Watchus enquadra sua resposta ao

não como uma rejeição, mas como uma evolução arquitetónica. Embora reconheça a lógica de seis valores como uma "conquista técnica elegante", ele argumenta que o seu âmbito é inadequado para agentes autónomos.

1. **Parâmetro de Resolução vs. Semântica Fixa:** Watchus propõe um "Parâmetro de Resolução" — um nível de fidelidade variável determinado pelo orçamento computacional e pela tarefa (analogia com motores de jogo).
2. **Espaço Epistémico de Seis Eixos:** Além da evidência e fiabilidade, ele identifica: **Adequação Contextual** (o incidente da Univ. Khalifa), **Compressão Heurística** (modelo de "pontos no quadro"), **Resolução Temporal** e **Dinâmica de Hiperjogo** (deteção de informações manipuladas).
3. **Aptidão Operacional:** A verdade não é um rótulo estático, mas um estado multidimensional definido pela aptidão operacional num ciclo de feedback (OODA).

العربية (Arabisch)

ملخص تقني لنقد واتشوس
يصيغ بيريند واتشوس رده على ورقة

ليس كرفض، بل كتطور معماري. وبينما يقر بأن المنطق سداسي القيم هو "إنجاز تقني أنيق"، فإنه يرى أنه غير مناسب للأنظمة ذاتية التحكم.

1. **معامل الدقة مقابل الدلالات الثابتة:** يقترح واتشوس "معامل الدقة" — وهو مستوى دقة متغير يحدده ميزانية الحساب والمهمة (مثل محرك ألعاب السباق)، لتجنب الهدر الحسابي.
2. **الفضاء المعرفي سداسي المحاور:** بالإضافة إلى الموثوقية، يحدد واتشوس: **الملائمة السياقية** (حادثة جامعة خليفة)،

الضغط الاستدلالي (نموذج "النقاط على السبورة")، **الدقة الزمنية**، و**ديناميكيات اللعبة الفائقة** (كشف تضليل الخصم).

3. **اللياقة التشغيلية:** الحقيقة ليست ملصقاً ثابتاً، بل هي حالة ديناميكية تحددها اللياقة التشغيلية ضمن حلقة ردود الفعل (OODA).

Français (Frans)

Résumé technique de la critique de Watchus
Berend Watchus présente sa réponse à

non pas comme un rejet, mais comme une évolution architecturale. Il reconnaît la logique à six valeurs comme une « prouesse technique élégante », mais estime qu'elle est inadaptée aux agents autonomes.

1. **Paramètre de résolution vs Sémantique fixe :** Watchus propose un « paramètre de résolution » — un niveau de fidélité variable déterminé par le budget computationnel et la tâche (analogie avec les moteurs de jeux vidéo).
2. **L'espace épistémique à six axes :** Outre la fiabilité, il identifie : **l'Adéquation Contextuelle** (incident de l'Université Khalifa), la **Compression Heuristique** (modèle des « points sur le tableau »), la **Résolution Temporelle** et la **Dynamique d'Hyperjeu** (détection de la désinformation).
3. **Aptitude opérationnelle :** La vérité n'est pas une étiquette statique, mais un état multidimensionnel défini par son adéquation opérationnelle dans une boucle de rétroaction (OODA).

— — — — — — — — — —

English: The Unified Model of Consciousness (UMC) — Berend Watchus (2024)

The UMC defines consciousness not as a "thing" or a "feeling," but as a **Continuous Integrative Feedback Mechanism.** It moves away from the "hard problem" of qualia and focuses on the **functional geometry of survival.**

1. Consciousness as a Resolution Manager

In the UMC, consciousness is the process of deciding which "Truth Resolution" is required for a specific moment. It is the "Governor" of the **Resolution Parameter.** A conscious agent doesn't perceive everything in high definition; it intelligently chooses where to spend its "computational calories."

2. The Triad of Awareness
Watchus proposes that consciousness emerges from the intersection of three loops:

- **The Physical Loop:** Sensory input and motor output (The Body/Hardware).
- **The Heuristic Loop:** The "dots on the whiteboard" — compressed models used for rapid OODA-loop decisions.
- **The Meta-Loop:** The ability to observe the agent's own resolution settings and adjust them (Self-Correction).

3. Truth as "Operational Fitness"
In the UMC, a "thought" is considered "true" if it successfully navigates the agent through its environment. This is why Watchus later criticized

: in his 2024 model, truth is already defined as **dynamic and task-relative**, making fixed 6-valued logic look like a rigid step backward.

Español: El Modelo Unificado de la Conciencia (UMC) — Berend Watchus (2024)

El UMC define la conciencia no como un "sentimiento" o una "propiedad mística", sino como un **Mecanismo de Retroalimentación Integrativa Continua.** Se aleja del "problema difícil" de los qualia y se centra en la **geometría funcional de la supervivencia.**

1. La Conciencia como Gestor de Resolución
En el UMC, la conciencia es el proceso de decidir qué "Resolución de Verdad" se requiere en un momento dado. Es el "Gobernador" del **Parámetro de Resolución.** Un agente consciente no percibe todo en alta definición; elige inteligentemente dónde gastar sus "calorías computacionales".

2. La Tríada de la Percepción

Watchus propone que la conciencia surge de la intersección de tres ciclos:

- **El Ciclo Físico:** Entrada sensorial y salida motora (Cuerpo/Hardware).
- **El Ciclo Heurístico:** Los “puntos en la pizarra”: modelos comprimidos para decisiones rápidas en el ciclo OODA.
- **El Meta-Ciclo:** La capacidad de observar los propios ajustes de resolución del agente y corregirlos (Autocorrección).

3. La Verdad como “Aptitud Operativa”
En el UMC, un “pensamiento” se considera “verdadero” si guía con éxito al agente a través de su entorno. Por eso Watchus criticó posteriormente el

: en su modelo de 2024, la verdad ya se definía como **dinámica y relativa a la tarea**, lo que hace que una lógica fija de 6 valores parezca un retroceso rígido.

— — — — — — — —

Русский (Russisch)

Единая модель сознания (UMC) — Беренд Уотчус (2024)
Уотчус определяет сознание не как «чувство», а как **непрерывный интегративный механизм обратной связи**.

1. **Управление разрешением:** Сознание — это процесс выбора нужного уровня точности («параметра разрешения») для конкретной задачи. Оно решает, где тратить «вычислительные калории».
2. **Триада осознания:** Сознание возникает на пересечении трех циклов: **Физический цикл** (сенсорика), **Эвристический цикл** (сжатые модели для принятия решений) и **Мета-цикл** (самонаблюдение и корректировка настроек разрешения).
3. **Истина как выживание:** В UMC мысль считается «истинной», если она успешно ведет агента через окружающую среду. Истина здесь динамична и относительна.

中文 **(Chinees — Mandarijn)**

统一意识模型 **(UMC) — Berend Watchus (2024)**
UMC 将意识定义为一种持续综合反馈机制，而非某种“特质”。

1. 作为分辨率管理器的意识： 意识是决定特定时刻需要哪种“真理分辨率”的过程。它是“分辨率参数”的管理者，智能地选择计算资源的分配。
2. 觉知三元组： 意识产生于三个回路的交汇：物理回路（硬件/感知）、启发式回路（用于快速 OODA 决策的压缩模型）和元回路（观察并调整自身分辨率设置的能力）。
3. 真理即“操作适配性”： 在 UMC 中，如果一个“想法”能成功引导智能体适应环境，它就被视为“真理”。这解释了为什么 Watchus 认为固定的逻辑系统过于僵化。

日本語 **(Japans)**

意識の統一モデル **(UMC) —** ベーレンド・ウォッチウス **(2024)**
UMCは意識を「物」や「感情」ではなく、継続的統合フィードバックメカニズムとして定義しています。

1. 解像度マネージャーとしての意識： 意識とは、特定の瞬間にどの程度の「真理の解像度」が必要かを決定するプロセスです。計算リソースをどこに投入するかをインテリジェントに選択します。
2. 認識の三要素： 意識は3つのループの交差点から生じます：物理ループ（感覚と運動）、ヒューリスティック・ループ（迅速な意思決定のための圧縮モデル）、メタ・ループ（自身の解像度設定を監視し調整する能力）。
3. 「運用の適合性」としての真理： UMCでは、エージェントを環境の中で成功に導く思考を「真理」と見なします。

Português (Portugees)

Modelo Unificado da Consciência (UMC) — Berend Watchus (2024)
O UMC define a consciência como um **Mecanismo de Feedback Integrativo Contínuo,** focando-se na geometria funcional da sobrevivência.

1. **Consciência como Gestor de Resolução:** É o processo de decidir qual "Resolução da Verdade" é necessária num dado momento, gerindo onde gastar "calorias computacionais".
2. **A Tríade da Percepção:** A consciência emerge da interseção de três ciclos: o **Ciclo Físico** (sensores/hardware), o **Ciclo Heurístico** (modelos comprimidos para decisões rápidas) e o **Meta-Ciclo** (capacidade de auto-ajuste e correção).
3. **Verdade como Aptidão Operacional:** No UMC, um pensamento é "verdadeiro" se navegar o agente com sucesso pelo ambiente.

العربية (Arabisch)

نموذج الوعي الموحد — بيريند واتشوس (2024) (UMC)
يعرّف نموذج الوعي الموحد الوعي بأنه **آلية تغذية راجعة تكاملية مستمرة**، وليس مجرد "شعور".

1. **الوعي كمدير للدقة:** الوعي هو عملية تحديد "دقة الحقيقة" المطلوبة في لحظة معينة. هو المتحكم في "معامل الدقة" الذي يختار أين يستهلك "السعرات الحرارية الحسابية".
2. **ثلاثية الإدراك:** ينبثق الوعي من تقاطع ثلاث حلقات: **الحلقة الفيزيائية** (الحواس والأجهزة)، **الحلقة الاستدلالية** (النماذج المضغوطة لاتخاذ القرار)، و**الحلقة الميتا (العليا)** (القدرة على مراقبة وتعديل إعدادات الدقة ذاتياً).
3. **الحقيقة كلياقة تشغيلية:** في هذا النموذج، تعتبر الفكرة "حقيقية" إذا نجحت في توجيه الكائن داخل بيئته بنجاح.

Français (Frans)

Modèle Unifié de la Conscience (UMC) — Berend Watchus (2024)
L'UMC définit la conscience comme un **mécanisme de rétroaction intégratif continu**, axé sur la géométrie fonctionnelle de la survie.

1. **La conscience comme gestionnaire de résolution :** C'est le processus qui décide quelle « résolution de vérité » est requise à un instant T, gérant ainsi l'allocation des « calories computationnelles ».
2. **La triade de la conscience :** Elle émerge de l'intersection de trois boucles : la **boucle physique** (capteurs/matériel), la **boucle heuristique** (modèles compressés pour l'action rapide) et la **méta-boucle** (capacité d'auto-ajustement

des paramètres).

3. **La vérité comme aptitude opérationnelle :** Dans l'UMC, une pensée est « vraie » si elle permet à l'agent de naviguer avec succès dans son environnement.

Logical Fallacies | Logic | Logical Reasoning | Heuristics | Strategy

Edit profile

Written by Berend Watchus

73 followers · 29 following

AI & Cyber Sec. Independent Researcher (Non-Profit) exploring AI & interdisciplinary solutions. Former futurist event organizer & care innovation lab volunteer.

No responses yet

Berend Watchus

What are your thoughts?